The Dynamic Logic of Permission

Ron van der Meyden

Computing Science,
University of Technology, Sydney,
PO Box 123, Broadway, NSW 2007,
Australia.

email: ron@socs.uts.edu.au

To appear in Journal of Logic and Computation

Abstract

Applications of deontic logic in areas such as intelligent legal information systems require the ability to represent at least two different notions of permission, one of which, "free choice permission," cannot be adequately represented in standard modal logics. We define a logic which handles this modality by using ideas from dynamic logic. The logic is also able to expresses a different notion of permission corresponding to a lack of prohibition. The main result is the completeness of an axiomatization of the logic.

1 Introduction

The study of the modalities of permission, obligation and prohibition is known in the philosophical literature as *deontic logic*. The deontic modalities are becoming increasingly of interest in computer science, with proposed applications including intelligent legal information systems [8], computer security [5, 15], software engineering [13, 14], database integrity constraints [10], and agent-oriented programming [21]. In this paper we consider the modality of permission, and present a logic in which a variety of modalities of permission may be expressed. Our work was motivated in the context of a program of research dealing with legal applications, but the presentation is abstract, and may well be of interest in other applications of deontic logic in computer science.

In the context of intelligent legal information systems, the relevance of the deontic modalities is not just that the law itself prescribes permissions, obligations and prohibitions, but that in many legal domains, such as taxation law and the law of contracts, statutes and cases are stated in terms of these modalities. For example, the system EPS-II proposed by Schlobohm and McCarty [19] aims to provide support for lawyers engaged in estate planning, a legal domain in which clients' goals typically involve the minimization of income tax. The relevant taxation law makes explicit use of deontic concepts. More significantly, the input to this system would include a statement of rules governing a trust, which is to be analyzed for tax consequences. These rules concern certain permissions, obligations and prohibitions of the trustor and the trustee. Thus it is not possible to contemplate constructing such an information system without providing some logical means by which these modalities can be stated.

At least two types of permission modalities occur in legal contexts, *free choice permission* and *lack of prohibition*. For example, an hypothetical trust considered by Schlobohm and McCarty contains the following language:

> During the Trustor's lifetime, the Trustee may apply for the benefit of the Trustor and the Trustor's issue (or any other person or entity) such amounts of income and principal of the trust estate and in such proportion as the Trustee in the Trustee's absolute discretion shall determine.

The permission of the trustee being stated here is of the free choice kind. The trustee may distribute the income in whatever way he chooses. That is, he may perform the abstract action "distribute income" using *any* of the concrete actions satisfying this description. On the other hand, Section 2038 of the Internal Revenue Code states that the value of the gross estate includes

> ...the value of all property...where the enjoyment thereof was subject at the date of his death to any change through the exercise of a power (in whatever capacity exercisable) by the decedent alone or by the decedent in conjunction with any other person...to alter, amend, revoke or terminate.

The notion of permission in this clause involves a lack of prohibition. The clause states that if there is *some* way in which the trustor was permitted to alter the the trust, then the property shall be counted in the gross value.

The notion of free choice permission is notorious as the source of one of the numerous "paradoxes" of deontic logic. Unlike the set-theoretic paradoxes, these are not real paradoxes, but rather inadequacies of standard modal logics for the expression of deontic modalities. For example, if we represent free choice permission by a possibility operator M in a normal modal logic, we obtain the following undesirable valid formula: $Mp \supset M(p \vee q)$, a variant of the Ross paradox [18]. To see that this formula is undesirable, consider the instance "If it is permitted to mail the letter then it is permitted to mail the letter or burn it." Since, if M is to be interpreted as free choice permission, we should have $M(p \vee q) \supset Mq$ valid, this would yield the absurd conclusion that if it is permitted to mail the letter then it is permitted to burn it.

Various philosophers [1, 22] have argued that the paradoxes of deontic logic will be resolved in deontic logics in which it is *practitions* or *actions* rather than *propositions* which are the objects of the deontic modalities. The introduction by Pratt [16] of dynamic logic, a logic designed to reason about programs, made available a framework on which to base such logics. There have been attempts to construct deontic logics based on dynamic logic, by McCarty [8] and Meyer [12]. McCarty's logic incorporates many features, which cause considerable semantic complexity. Although McCarty lists a number of valid formulae he provides no completeness result. In this paper we introduce a dynamic logic of permission, DLP, a relative of McCarty's deontic logic, which distills the ideas necessary for a treatment of free choice permission. Our main contribution is to provide a complete axiomatization of this simpler logic.

Dynamic logic has two sorts of expressions: *programs* (which we will refer to as *actions*) and *propositions*. Propositions *hold* of states of affairs, whereas actions produce a *change* of state. Actions may be nondeterministic, having different ways in which they can be executed. The primary logical construct of dynamic logic is the modality $\langle\alpha\rangle X$, expressed X *after* α. This modality operates on an action α and a proposition X, and asserts that some execution of the action α results in a state satisfying the proposition X. Dually, we may assert that all executions of the action result in a state of affairs satisfying a proposition.

The semantics of dynamic logic provides for each action an explicit denotation describing the ways the action changes states of affairs. There have been two formulations of the action semantics. In the relational formulation, each action denotes a binary relation on states, such that a pair of states is in the relation when, if the action is performed in the first, the second may be the resulting state of affairs. For *Process Logic* [17], which allows description of the states of affairs *during* the execution of an action, not just those at its endpoints, a generalization of the relational semantics to *sequences* of states is required.

For the semantics of permission in Meyer [12], the relational semantics of actions suffices. He assumes there exists a set of permitted states of affairs, and takes an action to be permitted just when it results in a permitted state. Thus Meyer's logic of permission is obtained straightforwardly from dynamic logic by adding a new propositional constant **Permitted** and defining the permission operator P by $P(\alpha) \equiv \langle\alpha\rangle\mathbf{Permitted}$.

While Meyer's logic may adequately express a notion of permission related to static database integrity constraints [10], we do not feel that this operator captures an appropriate sense of permission of actions for legal applications. For example, the following is valid in Meyer's logic:

$$\langle\alpha\rangle P(\beta) \supset P(\alpha;\beta).$$

In this sentence the expression $\alpha;\beta$ expresses sequential composition of the actions. As an example, consider the sentence "If after shooting the president one has the right to remain silent then one has the right to shoot the president and remain silent." The antecedent of this sentence is true, but the consequent is false.

The key to a treatment of permission which avoids this undesirable inference is to use the process semantics for actions. Rather than identify *states* as permitted, we may then classify *executions* of actions as either "good" or "bad." Two different modalities of permission then suggest themselves. We might regard an action as permitted if at least one of its executions is "good." Alternatively, we could say that the permitted actions are those all of whose executions are "good." The former corresponds to the notion of *not forbidden*, the latter to *free-choice permission*.

The *Dynamic Logic of Permission* (henceforth DLP) introduced below models both of these modalities, using two operators $\Diamond$ and π. In order to allow description of the effects of actions we also include the dynamic operator *after*. In constructing the proof theory of the deontic modalities it is convenient to generalize them to take an additional propositional argument describing the states resulting from the execution of the action. However DLP is not a *dyadic* logic in the sense conventional in the deontic logic literature. Permissions dependent on the state of affairs are still expressed in DLP using a conditional: $Q \supset \pi(\alpha, R)$, rather than including the condition Q in the scope of the modality.

Section 2 describes the syntax and the semantics of DLP, and Section 3 discusses related work and possible objections to this semantics. In Section 4 an axiomatization is described, which we prove complete in Section 5. Section 6 contains concluding remarks.

2 Syntax and Semantics

Let $\mathbf{Act} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \ldots\}$ be a set of *atomic action constants* and $\mathbf{Prop} = \{\mathrm{P}, \mathrm{Q}, \mathrm{R}, \ldots\}$ a set of *atomic propositions*. The set of *action expressions* is defined to be the smallest set $\mathbf{A}$ containing $\mathbf{Act}$ such that if $\alpha, \beta \in \mathbf{A}$ then $\alpha \cup \beta \in \mathbf{A}$, $\alpha;\beta \in \mathbf{A}$ and $\alpha^* \in \mathbf{A}$. The set of *propositions* of DLP is the smallest set containing $\mathbf{Prop}$ and satisfying:

1. If X and Y are propositions then $X \vee Y$ and $\neg X$ are propositions.

2. If α is an action expression and X is a proposition then $\langle\alpha\rangle X$, $\Diamond(\alpha, X)$ and $\pi(\alpha, X)$ are propositions.

We use the customary abbreviations from propositional logic, as well as the abbreviations $[\alpha]X$ for $\neg\langle\alpha\rangle\neg X$ and $\bigcirc(\alpha, X)$ for $\neg\Diamond(\alpha, \neg X)$. We will refer to the operators $\langle\rangle$, $\Diamond$ and π as *after, not-forbidden* and *permitted* respectively. The latter will correspond to the notion of free choice permission.

Formally, the model theory of DLP is constructed by augmenting the *process semantics* for Propositional Dynamic Logic [17], in which the denotation of an action expression is a set of sequences of world states. A sequence of states is in the denotation of an action just in case some way of performing the action results in that sequence of states. While the semantics of the modality *after* refers only to the initial and final states of these sequences, the semantics of the deontic operators refer to the sequences themselves. Whether or not an action is permitted depends not only on the world state resulting from the execution of the action, but also on the intermediate states.

If S is a set then S^+ denotes the set of all finite sequences of elements of S of length one or more. If σ is a finite sequence we will write $\sigma[i]$ for the first state of σ and $\sigma[f]$ for the final state. If $\sigma_1 = (s_1, \ldots, s_n)$ and $\sigma_2 = (t_1, \ldots, t_m)$ are two sequences with $s_n = t_1$ then $\sigma_1; \sigma_2$ is defined to be the sequence $(s_1, \ldots, s_n, t_2, \ldots, t_m)$. Otherwise $\sigma_1; \sigma_2$ is undefined. If T and U are sets of sequences of elements of a set S then

$$T; U = \{\sigma_1; \sigma_2 \mid \sigma_1 \in T,\ \sigma_2 \in U \text{ and } \sigma_1; \sigma_2 \text{ is defined}\}$$

and

$$T^* = \{(w) \mid w \in S\} \cup T \cup T; T \cup \ldots$$

A model for the language DLP is a tuple $M = \langle S, P, \tau, V\rangle$ where S is a set of states, P is a set of pairs of elements of S, the function $\tau : \mathbf{Act} \rightarrow \mathcal{P}(\mathrm{S}^+)$ maps atomic action symbols to sets of finite sequences and $V : \mathbf{Prop} \times \mathrm{S} \rightarrow \{\mathbf{true}, \mathbf{false}\}$ is a valuation on the states. Intuitively, P tells us which one-step state transitions are permitted. A sequence of state transitions $\sigma = (s_1, \ldots, s_n)$ is permitted just in case each transition (s_i, s_{i+1}) is in P. We will say in this case that the sequence is *green*. Note that this definition makes all sequences of length one green.

The mapping τ gives the interpretation of atomic actions. Intuitively, if $\mathbf{a}$ is an atomic action then $\tau(\mathbf{a})$ is the set of all sequences of states which constitute an execution of the action $\mathbf{a}$. The mapping τ is extended to give interpretations to all action expressions by

$$\begin{aligned} \tau(\alpha; \beta) &= \tau(\alpha); \tau(\beta) \\ \tau(\alpha \cup \beta) &= \tau(\alpha) \cup \tau(\beta) \\ \tau(\alpha^*) &= \tau(\alpha)^*. \end{aligned}$$

That is, the action $\alpha \cup \beta$ can be performed by executing either α or β. The action $\alpha; \beta$ consists of executing first α then β, and the action α^* is performed by executing α some finite number of times. If $w \in S$ is a state then we will write $\tau_w(\alpha)$ for the set of all sequences in $\tau(\alpha)$ with initial state w. In other words, $\tau_w(\alpha)$ is the set of all ways that α may be performed, starting in state w.

Satisfaction of sentences of DLP in a model $M = \langle S, P, \tau, V\rangle$ at a state $w \in S$ is defined as follows:

$$\begin{array}{ll} M, w \models Q & \text{if } V(Q, w) = \mathbf{true}, \text{ when } Q \in \mathbf{Prop} \\ M, w \models X \vee Y & \text{if } M, w \models X \text{ or } M, w \models Y \\ M, w \models \neg X & \text{if not } M, w \models X \\ M, w \models \langle\alpha\rangle\, X & \text{if } M, \sigma[f] \models X \text{ for some } \sigma \in \tau_w(\alpha) \\ M, w \models \Diamond(\alpha, X) & \text{if } M, \sigma[f] \models X \text{ for some green } \sigma \in \tau_w(\alpha) \\ M, w \models \pi(\alpha, X) & \text{if all } \sigma \in \tau_w(\alpha) \text{ with } M, \sigma[f] \models X \text{ are green.} \end{array}$$

A sentence X is valid if $M, w \models X$ for all models M and all states w of M.

Intuitively, $\langle\alpha\rangle X$ is to be thought of as asserting that, starting from the current state, some execution of the action α ends in a state in which the proposition X holds. The proposition $\Diamond(\alpha, X)$ asserts that X holds after some execution of α which is not forbidden, i.e. in which every step is legal. Finally, $\pi(\alpha, X)$ expresses free choice permission: it asserts that each execution of α which ends in a state satisfying X is permitted. That is, provided α is performed (from the current state) in such a way as to terminate in a state satisfying X, no prohibited state transition will be made.

The dual $[\alpha]X$ of *after* asserts that *all* executions of α result in a state satisfying X. The dual of *not-forbidden* is a weak form of obligation: $\bigcirc(\alpha, X)$ asserts that all permitted executions of α result in a state satisfying X. In other words, it is obligatory, if performing α, to perform it in such a way that X becomes true. (It does not follow that α is not forbidden, nor that other actions are forbidden.) Monadic deontic operators may be obtained by setting the propositional argument to be **true**. Thus, we write $\Diamond(\alpha)$ for $\Diamond(\alpha, \mathbf{true})$, which holds when some execution of α is permitted, and we write $\pi(\alpha)$ for $\pi(\alpha, \mathbf{true})$, which holds when all executions of α are permitted.

3 Discussion

We now show that the permission operators we have defined avoid the paradoxes discussed in the introduction, and consider a number of possible objections to our proposed semantics. We note that our primary concern in this paper is to develop a proof theory for the interaction of action modalities and the particular notion of free choice permission identified in the work of McCarty [8]. The semantics of the operators defined above were selected for their simplicity and their ability to provide for this proof theory a parsimonious basis, and are not intended to have a direct correspondence to natural language constructions. We believe that our definitions are natural, and sufficiently expressive to allow for the formalization and disambiguation of a range of natural language constructions involving permission, but this is an empirical question which we will not pursue in depth. We do not wish to claim that the operators we have defined exhaust the range of naturally occurring permission expressions.

To see that operators we have defined avoid the paradoxes mentioned in the introduction, let us first consider the formula $\langle\alpha\rangle P(\beta) \supset P(\alpha;\beta)$ we found problematical in Meyer's logic. There are two corresponding formulae in our notation:

$$\langle\mathbf{a}\rangle\pi(\mathbf{b}) \supset \pi(\mathbf{a};\mathbf{b}) \tag{1}$$

when we interpret the modality P as free choice permission, and

$$\langle\mathbf{a}\rangle\Diamond(\mathbf{b}) \supset \Diamond(\mathbf{a};\mathbf{b}) \tag{2}$$

when we interpret P as lack of prohibition. Take M to be the model with states u, v and w and let the interpretations of the primitive actions be given by $\tau(\mathbf{a}) = \{(u,v)\}$ and $\tau(\mathbf{b}) = \{(v,w)\}$. If $(v,w) \in P$ is a permitted transition, but $(u,v) \notin P$ is not a permitted transition, then the sequence (v,w) is green, but neither the sequence (u,v) nor $(u,v);(v,w) = (u,v,w)$ is green. Since $\tau(\mathbf{a};\mathbf{b}) = \{(u,v,w)\}$, this means that both the formulae (1) and (2) are false at the state u, since their antecedents are true but their consequents are false.

Next, the paradox of free choice permission, $Mp \supset M(p \vee q)$ is not valid for the free choice permission operator π.[1] That is, expressed in our notation, the formula

[1] We note that the formula $\Diamond(\alpha) \supset \Diamond(\alpha\cup\beta)$ is valid in our logic. Given our intuitive interpretation of $\Diamond$ as expressing a lack of prohibition, this is as it should be.

$\pi(\alpha) \supset \pi(\alpha \cup \beta)$ is not valid. This may be seen by constructing a model M with a state w in which $\tau_w(\alpha)$ contains a single sequence σ_1, which is green, and in which $\tau_w(\beta)$ contains a distinct sequence σ_2, which is not green. Because all sequences in $\tau_w(\alpha)$ are green, we then have $M, w \models \pi(\alpha)$. However, since $\tau_w(\alpha \cup \beta)$ contains σ_2, we do not have that $M, w \models \pi(\alpha \cup \beta)$. This resolution of the paradox of free choice permission is very closely related to McCarty's [8]: our semantics for permission may be viewed as a simplification if his more elaborate (and intuitionistic) framework.

On the other hand, the formula

$$\pi(\alpha \cup \beta) \equiv \pi(\alpha) \wedge \pi(\beta) \tag{3}$$

is easily seen to be valid. This expresses the intuitively correct claim that "it is permitted to do α or β" is equivalent to "it is permitted to do α and it is permitted to do β". It might be objected that this claim can be interpreted so as to be true, but without "it is permitted to do α" being interpreted as stating that every way of doing α is permitted. That is, what is permitted is to make the choice, but both actions may be subject to normative constraints. The formula $\pi(\alpha \cup \beta)$ is too strong to capture this sense of disjunctive permission, and the formula $\Diamond(\alpha \cup \beta)$ is too weak.

Indeed, Meyer [11] has proposed a different dynamic logic based resolution of the paradox of free choice permission, that also satisfies the formula (3), but without having the property that a permitted action is permissible in all its executions. Rather than deal with a new operator for permission, he introduces a new action construct $\oplus$ such that $\alpha \oplus \beta$ denotes an agent's "*internally* choosing between α and β", and distinguishes this from $\alpha + \beta$ (essentially equivalent to $\alpha \cup \beta$, in our notation) which represents an *external* choice, made by the environment. The semantics of the new operator involves a complication of the standard dynamic logic semantics that preserves aspects of an action expression's syntax in its semantics.[2] Using the reduction of deontic logic to dynamic logic of [12], Meyer then obtains the valid formula

$$P(\alpha \oplus \beta) \equiv P(\alpha) \wedge P(\beta).$$

Since it is based on the transition semantics for dynamic logic, this approach to free choice still suffers from the sequential "paradox". However, we believe the ideas it uses may be applied also to a process based semantics.

Given that there exist at least these two approaches to free choice permission solving the paradox while satisfying the formula (3), which is the right one? We would argue that both have their applications. Examples like

> It is permitted to ship low grade radioactive materials by sea or air, but it is not permitted to ship them to countries on the proscribed list. (A)

suggest a semantics like Meyer's. On the other hand, this kind of statement of permission does not provide sufficient information to an agent who must act. If an agent is to act on a permission to perform an abstract action α, that action must be be permissible in all its executions if the agent is to have confidence that no norm will be violated. An example of a statement providing this degree of permission is

[2] Only the operators $\oplus$ and $+$ are discussed in [11]. It is not made clear how the semantics of dynamic logic is to be generalized to handle action expressions involving the other dynamic operators, such as $(\alpha \oplus \beta);\gamma$.

> A client may withdraw from his savings account any amount, provided he maintains a positive balance.

This sort of statement is best modelled semantics using the free choice operator π studied in this paper, by means of a formula like $\pi(\mathbf{withdraw},\ \mathit{balance} > 0)$, where the intended semantics of **withdraw** is all sequences of events constituting the withdrawal of some amount.

Moreover, it seems that something like the universally quantified semantics of the operator π lurks beneath even statements like (A). It is natural to interpret a body of such statements as being maximally permissive to the extent consistent with the explicitly stated prohibitions. For example, if (A) is the only rule concerning the shipping of radioactive material, and if Italy is not on the proscribed list, then it is reasonable to infer that it is permitted to ship low grade material by air to Rome on December 25. Note however, that if the Italian government passes legislation forbidding the transport of radioactive material on holidays, then this inference from (A) would be blocked. Clearly, to capture this sort of phenomenon would require an excursion into the territory of non-monotonic logic. While we consider non-monotonicity in deontic logic to be an important area to be explored, we will not pursue this here.

We consider one final possible objection to the operator π. It might be argued that if an action is free choice permitted, then it must be possible. The operator π does not have this property, since if the semantics $\tau(\mathbf{a})$ of an action $\mathbf{a}$ is the empty set, then $\pi(\mathbf{a})$ holds at every state. We do not view this to be a serious problem. If a notion of free choice permission is required with the property that free choice permitted actions are possible, one can be defined by means of $P(\alpha, X) \equiv \pi(\alpha, X) \wedge \langle\alpha\rangle X$.

4 Axiomatization

We now introduce the axiomatization we will show to be complete for DLP. To handle the dynamic modality *after* we use the Segerberg [20] axiomatization of dynamic logic, which includes all propositional tautologies, plus the schemata

$$\begin{array}{ll}
\langle\rangle 1. & \langle\alpha\rangle\mathbf{false} \equiv \mathbf{false} \\
\langle\rangle 2. & \langle\alpha;\beta\rangle X \equiv \langle\alpha\rangle\langle\beta\rangle X \\
\langle\rangle 3. & \langle\alpha\cup\beta\rangle X \equiv \langle\alpha\rangle X \vee \langle\beta\rangle X \\
\langle\rangle 4. & \langle\alpha^*\rangle X \equiv X \vee \langle\alpha;\alpha^*\rangle X \\
\langle\rangle 5. & \langle\alpha\rangle(X \vee Y) \equiv \langle\alpha\rangle X \vee \langle\alpha\rangle Y \\
\langle\rangle 6. & X \wedge [\alpha^*](X \supset [\alpha]X) \supset [\alpha^*]X.
\end{array}$$

Axioms $\langle\rangle 2 - \langle\rangle 5$ enable complex action expressions to be "eliminated" from formulae. For the union and composition operators this is clear. For the iteration operator note that the combination of Axioms $\langle\rangle 2 - \langle\rangle 5$ suggests the "equivalence" $\langle\alpha^*\rangle X \equiv X \vee \langle\alpha\rangle X \vee \langle\alpha;\alpha\rangle X \vee \ldots$, which of course cannot be expressed in the logic. Axiom $\langle\rangle 6$, called the *induction* axiom, is used to avoid the need to manipulate such infinite disjunctions. Intuitively, the first occurrence of X in this axiom is the base case, and $X \supset [\alpha]X$ corresponds to the inductive step, universally quantified by the operator '$[\alpha^*]$.'

An examination of the semantics of the modality *not-forbidden* reveals that it is identical to the semantics of *after,* except that the quantification is over the green sequences in $\tau(\alpha)$ rather than all sequences in this set. Furthermore, restricting to green sequences preserves the equations defining the function τ on complex action expressions. For example, the set of green sequences in $\tau(\alpha \cup \beta)$ is equal to the union of the set of green sequences in $\tau(\alpha)$ with the set of green sequences in $\tau(\beta)$. This suggests that the internal logic of the modality *not-forbidden* is precisely the same as that of the operator *after*. Thus, we add the schemata

$$\begin{array}{ll}
\Diamond 1. & \Diamond(\alpha, X) \supset \langle\alpha\rangle X \\
\Diamond 2. & \Diamond(\alpha;\beta, X) \equiv \Diamond(\alpha, \Diamond(\beta, X)) \\
\Diamond 3. & \Diamond(\alpha \cup \beta, X) \equiv \Diamond(\alpha, X) \vee \Diamond(\beta, X) \\
\Diamond 4. & \Diamond(\alpha^*, X) \equiv X \vee \Diamond(\alpha;\alpha^*, X) \\
\Diamond 5. & \Diamond(\alpha, X \vee Y) \equiv \Diamond(\alpha, X) \vee \Diamond(\alpha, Y) \\
\Diamond 6. & X \wedge \bigcirc(\alpha^*, X \supset \bigcirc(\alpha, X)) \supset \bigcirc(\alpha^*, X).
\end{array}$$

Note that these axioms correspond directly to the respective axioms for *after*, except that $\Diamond 1$ and $\langle\rangle 1$ do not match. However, $\Diamond 1$ together with the axiom $\langle\rangle 1$ yields $\Diamond(\alpha, \mathbf{false}) \equiv \mathbf{false}$, so the logics of *not-forbidden* and *after* are indeed the same.

The axioms for *permitted* also mirror to some extent those of *after*:

$$\begin{array}{ll}
\pi 1. & [\alpha]\neg X \supset \pi(\alpha, X) \\
\pi 2. & \pi(\alpha;\beta, X) \equiv \pi(\alpha, \langle\beta\rangle X) \wedge [\alpha]\pi(\beta, X) \\
\pi 3. & \pi(\alpha \cup \beta, X) \equiv \pi(\alpha, X) \wedge \pi(\beta, X) \\
\pi 4. & \pi(\alpha^*, X) \equiv \pi(\alpha;\alpha^*, X) \\
\pi 5. & \pi(\alpha, X \vee Y) \equiv \pi(\alpha, X) \wedge \pi(\alpha, Y) \\
\pi 6. & [\alpha^*]\pi(\alpha, \langle\alpha^*\rangle X) \supset \pi(\alpha^*, X).
\end{array}$$

Axiom $\pi 3$ states that the operator π has the defining property of free choice permission. The axiom $\pi 1$ indicates that if there is no way to perform an action, then this action is free choice permitted. The axiom $\pi 6$ may be viewed as a kind of induction axiom, which enables the introduction of the action repetition operator within the scope of the permission operator.

Finally, we have some axioms that combine the three modalities:

$$\begin{array}{ll}
C1. & \pi(\alpha, X) \wedge \langle\alpha\rangle X \supset \Diamond(\alpha, X) \\
C2. & \Diamond(\alpha, X) \wedge [\alpha](X \supset Y) \supset \Diamond(\alpha, Y) \\
C3. & \pi(\alpha, Y) \wedge [\alpha](X \supset Y) \supset \pi(\alpha, X).
\end{array}$$

Note that π and $\Diamond$ operate in a dual fashion. This duality is also apparent when one compares $\Diamond 3$ with $\pi 3$ and $\Diamond 5$ with $\pi 5$. The rules of inference are *modus ponens* and *generalization*:

$$X, X \supset Y \vdash Y \qquad\qquad X \vdash [\alpha]X.$$

Observe that we do not need separate generalization rules for the other modalities because the derived rules

$$X \vdash \bigcirc(\alpha, X) \qquad\qquad \neg X \vdash \pi(\alpha, X)$$

follow from generalization using $\Diamond 1$ and $\pi 1$.

5 Soundness and Completeness

We will now establish the completeness of the axioms by modifying the proof of completeness of the Segerberg axioms for Propositional Dynamic Logic of Kozen and Parikh [6]. The verification of soundness is straightforward, and left to the reader. The proof of completeness uses a generalization of the Fisher-Ladner closure [4] of a proposition. The *deontic Fisher-Ladner closure* $Cl(W)$ of a proposition W of DLP is the smallest set of propositions containing W and satisfying the definition of the Fisher-Ladner closure:

1. if $X \vee Y \in Cl(W)$ then $X \in Cl(W)$ and $Y \in Cl(W)$,
2. if $\neg X \in Cl(W)$ then $X \in Cl(W)$,
3. if $\langle\alpha\rangle X \in Cl(W)$ then $X \in Cl(W)$,
4. if $\langle\alpha \cup \beta\rangle X \in Cl(W)$ then $\langle\alpha\rangle X \in Cl(W)$ and $\langle\beta\rangle X \in Cl(W)$,
5. if $\langle\alpha;\beta\rangle X \in Cl(W)$ then $\langle\alpha\rangle\langle\beta\rangle X \in Cl(W)$,
6. if $\langle\alpha^*\rangle X \in Cl(W)$ then $\langle\alpha\rangle X \in Cl(W)$ and $\langle\alpha\rangle\langle\alpha^*\rangle X \in Cl(W)$.

but in addition satisfying the rules:

7. if $\Diamond(\alpha, X) \in Cl(W)$ then $\langle\alpha\rangle X \in Cl(W)$,
8. if $\Diamond(\alpha \cup \beta, X) \in Cl(W)$ then $\Diamond(\alpha, X) \in Cl(W)$ and $\Diamond(\beta, X) \in Cl(W)$,
9. if $\Diamond(\alpha;\beta, X) \in Cl(W)$ then $\Diamond(\alpha, \Diamond(\beta, X)) \in Cl(W)$,
10. if $\Diamond(\alpha^*, X) \in Cl(W)$ then $\Diamond(\alpha, X) \in Cl(W)$ and $\Diamond(\alpha, \Diamond(\alpha^*, X)) \in Cl(W)$,
11. if $\pi(\alpha, X) \in Cl(W)$ then $\langle\alpha\rangle X \in Cl(W)$,
12. if $\pi(\alpha \cup \beta, X) \in Cl(W)$ then $\pi(\alpha, X) \in Cl(W)$ and $\pi(\beta, X) \in Cl(W)$,
13. if $\pi(\alpha;\beta, X) \in Cl(W)$ then $\pi(\alpha, \langle\beta\rangle X) \in Cl(W)$ and $[\alpha]\pi(\beta, X) \in Cl(W)$,
14. if $\pi(\alpha^*, X) \in Cl(W)$ then $\pi(\alpha;\alpha^*, X) \in Cl(W)$ and $[\alpha^*]\pi(\alpha, \langle\alpha^*\rangle X) \in Cl(W)$.

We will in the sequel refer to the deontic Fisher-Ladner closure as simply the *closure*. Note the correspondence between the axioms and the closure rules. Rules 8 - 10 for the modality *not-forbidden* correspond directly to the standard rules 4 - 6 for the modality *after*. Like the standard Fisher-Ladner closure, the deontic Fisher-Ladner closure $Cl(W)$ is finite for any W. The number of formulae in the standard closure is linear in the size of the formula. This implies that the number of formulae which can be generated by means of the rules 1-12 from a formula W is linear in the size of W. Note that Rule 13 decreases the complexity of the action expressions governed by the modality π. The formula $\pi(\alpha;\alpha^*, X)$ introduced by rule 14 matches the precondition only of rule 13, generating the further formulae $\pi(\alpha, \langle\alpha^*\rangle X)$ and $[\alpha]\pi(\alpha^*, X)$. Note that in the former we have reduced the complexity of the action expression governed by π, whereas the latter, by the standard closure rules, produces only a linear number of new formulae. A similar argument applies to the formula $[\alpha^*]\pi(\alpha, \langle\alpha^*\rangle X)$ introduced by rule 14.

From axiom $\langle\rangle 5$ we obtain the consequence $\langle\alpha\rangle X \wedge [\alpha](X \supset Y) \supset \langle\alpha\rangle Y$. Together with C2, C3 this yields via a simple induction:

Lemma 1 If $\vdash X \equiv Y$ then $\vdash W(X) \equiv W(Y)$ where $W(Y)$ is the result of substituting Y for an occurrence of X in $W(X)$.

We will say that a formula X is *consistent* if not $\vdash \neg X$. Note that if a disjunction $X_1 \vee \ldots \vee X_n$ is consistent then one of the X_i is consistent. If $\vdash X \supset Y$ then we will write $X \leq Y$. Suppose that W is a consistent formula. We will show that there exists a model $M = \langle S, P, \tau, V\rangle$ and state $w \in S$ such that $M, w \models W$. Let $Cl(W) = \{X_1, \ldots, X_n\}$. We will call an *atom* any consistent conjunction $A = Y_1 \wedge \ldots \wedge Y_n$ where each Y_i is either X_i or $\neg X_i$. If X is a proposition then we will write $X \in A$ to mean that X is one of the Y_i. There is at least one atom $A \leq W$, since any $X \in Cl(W)$ can be proved to be equivalent to the disjunction of all atoms $B \leq X$.

The model we construct will have states consisting of pairs (A, c) where A is an atom and c is an element of the set $\{\mathbf{red}, \mathbf{green}\}$. The denotation $\tau(\mathbf{a})$ of the atomic action symbol $\mathbf{a}$ will be the set of sequences of states of length two $(A_1, c_1), (A_2, c_2)$ such that

1. $A_1 \wedge \langle\mathbf{a}\rangle A_2$ is consistent,
2. If $\pi(\mathbf{a}, X) \in A_1$ for some $X \geq A_2$ then $c_2 = \mathbf{green}$,
3. If $\neg\Diamond(\mathbf{a}, X) \in A_1$ for some $X \geq A_2$ then $c_2 = \mathbf{red}$.

The permitted transitions will be those transitions in τ to states marked **green**:

$$P = \bigcup_{\mathbf{a} \in \mathbf{Act}} \{(A_1, c_1)(A_2, c_2) \in \tau(\mathbf{a}) \mid c_2 = \mathbf{green}\}.$$

Finally, the valuation V is defined by

$$V(Q, (A, c)) = \mathbf{true} \ \text{ iff } \ Q \in A.$$

Observe that the colour c of a state (A, c) affects the colour of arcs entering the state, but has no effect on the arcs leaving the state, so there exists an isomorphism between the submodels generated by $(A, \mathbf{red})$ and $(A, \mathbf{green})$. Consequently, these states support exactly the same formulae. We will refer to a state of the form (A, c) as an A-state. If $\sigma \in \tau(\alpha)$ is a sequence with initial state an A-state and final state a B-state, then we will write $\sigma : A \rightarrow^\alpha B$, or simply $A \rightarrow^\alpha B$ if we merely wish to assert the existence of such a sequence. If there exists a green sequence of this form we will write $A \Rightarrow^\alpha B$. Note that by the previous remark, if either of these relations holds then a sequence of the required sort exists from *any* A-state to *some* B-state. The following two lemmas are straightforward generalizations of results in [6] to our slightly more complex model:

Lemma 2 For any program α,

1. if $A \wedge \langle\alpha\rangle B$ is consistent then $A \rightarrow^\alpha B$,

2. if $A \wedge \Diamond(\alpha, B)$ is consistent then $A \Rightarrow^\alpha B$.

Proof The proof is by induction on the complexity of α. We will prove only the basis of part 1, since the remainder of the proof of part 1 is similar to the proof of part 2 below.

To show part 1 in the case $\alpha = \mathbf{a}$, we need to show that if $A \wedge \langle\alpha\rangle B$ is consistent, then there do not exist $X_1, X_2 \geq B$ such that $A \leq \pi(\alpha, X_1) \wedge \neg\Diamond(\alpha, X_2)$. Suppose that this is so. Then by generalization, C2 and C3, $A \leq \pi(\alpha, B) \wedge \neg\Diamond(\alpha, B)$, so by C1, $A \leq \neg\langle\alpha\rangle B$, contrary to assumption. It follows that there is an **a**-arc from any A-state to some B-state.

For the basis of part 2, if $A \wedge \Diamond(\alpha, B)$ is consistent, then by $\Diamond 1$, $A \wedge \langle\alpha\rangle B$ is consistent, and by C1, it is not the case that $A \leq \pi(\alpha, B)$. It follows by C3 that $(A, c)(B, \mathbf{green}) \in \tau(\alpha)$.

We now show the inductive step for part 2. We consider first the case $\alpha = \beta \cup \gamma$. If $A \wedge \Diamond(\beta \cup \gamma, B)$ is consistent, then by $\Diamond 3$ we also have that $(A \wedge \Diamond(\beta, B)) \vee (A \wedge \Diamond(\gamma, B))$ is consistent, so one of the disjuncts, say $(A \wedge \Diamond(\beta, B))$, is consistent. By the induction hypothesis, $A \Rightarrow^\beta B$, so $A \Rightarrow^{\beta\cup\gamma} B$.

Consider next the case $\alpha = \beta;\gamma$. If $A \wedge \Diamond(\beta;\gamma, B)$ is consistent, then by $\Diamond 2$ also $A \wedge \Diamond(\beta, \Diamond(\gamma, B))$ is consistent. Since **true** is provably equivalent to the disjunction of all atoms, by Lemma 1 and $\Diamond 5$ this implies that

$$\bigvee_C (A \wedge \Diamond(\beta, C \wedge \Diamond(\gamma, B)))$$

is consistent, where C varies over all atoms. Consequently $A \wedge \Diamond(\beta, C \wedge \Diamond(\gamma, B))$ is consistent for some C. By C2, $A \wedge \Diamond(\beta, C)$ is consistent, and by $\Diamond 1$ and $\langle\rangle 1$, $C \wedge \Diamond(\gamma, B)$ is consistent. It follows by the induction hypothesis that $A \Rightarrow^\beta C \Rightarrow^\gamma B$, so $A \Rightarrow^{\beta;\gamma} B$.

For the case $\alpha = \beta^*$, assume $A \wedge \Diamond(\beta^*, B)$ is consistent. Let Γ be the smallest set of atoms containing A such that if $C \in \Gamma$ and $C \wedge \Diamond(\beta, D)$ is consistent then $D \in \Gamma$. By the induction hypothesis $A \Rightarrow^{\beta^*} C$ for all $C \in \Gamma$, so it suffices to show $B \in \Gamma$. If $Y = \bigvee_{C \in \Gamma} C$ then by $\Diamond 5$ we have that

$$Y \wedge \Diamond(\beta, \neg Y) \equiv \bigvee_{C \in \Gamma, D \notin \Gamma} C \wedge \Diamond(\beta, D) \ .$$

But each of the disjuncts on the right is inconsistent, so $Y \wedge \Diamond(\beta, \neg Y)$ is also inconsistent. Thus $\vdash Y \supset \neg\Diamond(\beta, \neg Y)$ so by generalization and C2, $\vdash \neg\Diamond(\beta^*, \neg(Y \supset \neg\Diamond(\beta, \neg Y)))$ and consequently $\vdash A \supset \neg\Diamond(\beta^*, \neg(Y \supset \neg\Diamond(\beta, \neg Y)))$. Since also $\vdash A \supset Y$, by $\Diamond 6$, $\vdash A \supset \neg\Diamond(\beta^*, \neg Y)$, so $A \wedge \Diamond(\beta^*, \neg Y)$ is inconsistent. Were $B \notin \Gamma$ we would have $B \leq \neg Y$ and by generalization and C2, $A \wedge \Diamond(\beta^*, B)$ inconsistent, contrary to assumption. This establishes that $B \in \Gamma$.

Lemma 3 1. For any $\langle\alpha\rangle X \in Cl(W)$ and atom A, $A \leq \langle\alpha\rangle X$ iff $A \rightarrow^\alpha B$ for some $B \leq X$.

2. For any $\Diamond(\alpha, X) \in Cl(W)$ and atom A, $A \leq \Diamond(\alpha, X)$ iff $A \Rightarrow^\alpha B$ for some $B \leq X$.

Proof Again we establish only part 2. The proof of part 1 follows that of part 2 closely. For the implication from left to right, observe that if $A \leq \Diamond(\alpha, X)$ then by $\Diamond 5$, $A \wedge \Diamond(\alpha, B)$ is consistent for some $B \leq X$. By Lemma2, $A \Rightarrow^{\alpha} B$.

The proof from right to left is by induction on the complexity of α. For the basis, suppose that $A \Rightarrow^{\alpha} B$ and $B \leq X$, where $\alpha = \mathbf{a}$. Then, by the definition of the model, $\neg\Diamond(\alpha, Y) \in A$ for no $Y \geq B$. In particular, putting $Y = X$ we have $A \leq \Diamond(\alpha, X)$.

Consider now the case $\alpha = \beta \cup \gamma$. If $A \Rightarrow^{\beta\cup\gamma} B \leq X$ then $A \Rightarrow^{\beta} B$, say. By the induction hypothesis, $A \leq \Diamond(\beta, X)$, since by closure rule 8, $\Diamond(\beta, X) \in Cl(W)$. Now by $\Diamond 3$ it follows that $A \leq \Diamond(\beta \cup \gamma, X)$.

Next, consider the case $\alpha = \beta;\gamma$. If $A \Rightarrow^{\beta;\gamma} B \leq X$ then for some atom C we have $A \Rightarrow^{\beta} C \Rightarrow^{\gamma} B$. By closure rules 9, 7 and 3, $\Diamond(\gamma, X) \in Cl(W)$ so by the induction hypothesis $C \leq \Diamond(\gamma, X)$. By closure rule 9, $\Diamond(\beta, \Diamond(\gamma, X)) \in Cl(W)$ so by the induction hypothesis $A \leq \Diamond(\beta, \Diamond(\gamma, X))$. It follows by $\Diamond 2$ that $A \leq \Diamond(\beta;\gamma, X)$.

Finally, consider the case $\alpha = \beta^*$. If $A \Rightarrow^{\beta^*} B \leq X$ then there exist $A_1, \ldots, A_n$ such that $A = A_1 \Rightarrow^{\beta} A_2 \Rightarrow^{\beta} \ldots \Rightarrow^{\beta} A_n = B$. By $\Diamond 4$ we have that $A_n \leq \Diamond(\beta^*, X)$. Since $\Diamond(\beta, \Diamond(\beta^*, X)) \in Cl(W)$ by closure rule 10, $A_{n-1} \leq \Diamond(\beta, \Diamond(\beta^*, X))$, so by $\Diamond 2$ and $\Diamond 4$, $A_{n-1} \leq \Diamond(\beta^*, X)$. Repeating this argument yields $A \leq \Diamond(\beta^*, X)$.

We now give two lemmas which establish an equivalence similar to those in Lemma 3 for the remaining operator π.

Lemma 4 For any $\pi(\alpha, X) \in Cl(W)$ and atom A, if $A \leq \pi(\alpha, X)$ then for every A-state w, all sequences in $\tau_w(\alpha)$ ending in a state with atom $B \leq X$ are green.

Proof By induction on the complexity of α. For $\alpha = \mathbf{a}$ this is immediate from the construction of the model.

Consider the case $\alpha = \beta \cup \gamma$. Let $A \leq \pi(\beta \cup \gamma, X) \in Cl(W)$ and $\sigma \in \tau_w(\beta \cup \gamma)$, say $\sigma \in \tau_w(\beta)$. By $\pi 3$, $A \leq \pi(\beta, X)$. Since by closure rule 12, $\pi(\beta, X) \in Cl(W)$, the induction hypothesis implies that σ is green.

Consider next the case $\alpha = \beta;\gamma$. If $A \leq \pi(\beta;\gamma, X) \in Cl(W)$ and $\sigma \in \tau_w(\beta;\gamma)$ then $\sigma = \sigma_1;\sigma_2$ for some $\sigma_1 \in \tau_w(\beta)$ and $\sigma_2 \in \tau_{\sigma_1[f]}(\gamma)$ where $\sigma_1[f] = (C, c)$ for some atom C. Since $\langle\gamma\rangle X \in Cl(W)$ by closure rules 13 and 11, we have that $C \leq \langle\gamma\rangle X$ by Lemma 3. By $\pi 2$, $A \leq \pi(\beta, \langle\gamma\rangle X)$ and since $\pi(\beta, \langle\gamma\rangle X) \in Cl(W)$ by closure rule 13, it follows that σ_1 is green by the induction hypothesis. Also by $\pi 2$ we have that $A \leq [\beta]\pi(\gamma, X)$. Since $[\beta]\pi(\gamma, X) \in Cl(W)$ by closure rule 13, the contrapositive of the equivalence in Lemma 3(1) implies that $C \leq \pi(\gamma, X)$. By the inductive hypothesis, it follows that σ_2 is green. Thus $\sigma = \sigma_1;\sigma_2$ is green.

Finally, consider the case $\alpha = \beta^*$. Suppose $A \leq \pi(\beta^*, X)$ and that $\sigma = \sigma_0;\ldots;\sigma_n$ is a sequence of states with $\sigma_i \in \tau(\beta)$ and $\sigma[f] \leq X$. Let A_i be the initial state of σ_i and let $A_{n+1} = \sigma[f]$. By $\pi 4$ we have $A_0 \leq \pi(\beta;\beta^*, X)$ so by $\pi 2$, $A_0 \leq \pi(\beta, \langle\beta^*\rangle X)$ and $A_0 \leq [\beta]\pi(\beta^*, X)$. By Lemma 3 we have that $A_1 \leq \langle\beta^*\rangle X$ since $\langle\beta^*\rangle X \in Cl(W)$ by closure rule 11. It follows that σ_0 is green by the induction hypothesis. Also by Lemma 3, $A_1 \leq \pi(\beta^*, X)$, since $[\beta]\pi(\beta^*, X) \in Cl(W)$ by closure rules 14 and 13. Repeating this argument yields that σ is green.

Lemma 5 For any $\pi(\alpha, X) \in Cl(W)$ and atom A, if for every A-state w all sequences in $\tau_w(\alpha)$ ending in a state $B \leq X$ are green then $A \leq \pi(\alpha, X)$.

Proof Suppose that for all A-states w all sequences in $\tau_w(\alpha)$ ending in a state $B \leq X$ are green. We show by induction on the complexity of α that $A \leq \langle \pi(\alpha, X)$.

If $\alpha = \mathbf{a}$ then the sequences $\tau_w(\alpha)$ are of the form $(A, c_1)(B, c_2)$. By construction of the model we have that if all such sequences ending in a state $B \leq X$ are green then for all $B \leq X$ such that $A \wedge \langle \mathbf{a} \rangle B$ is consistent there exists an $X_1 \geq B$ with $\pi(\mathbf{a}, X_1) \in A$. By C3 it follows that $A \leq \pi(\mathbf{a}, B)$ for such B. If $B \leq X$ is such that $A \wedge \langle \alpha \rangle B$ is inconsistent then $A \leq [\alpha] \neg B$ so by $\pi 1$ we also have $A \leq \pi(\mathbf{a}, B)$ for the remaining $B \leq X$. Since $X \in Cl(W)$ by closure rules 11 and 3, X is provably equivalent to the disjunction of all atoms $B \leq X$. It now follows by $\pi 5$ and Lemma 1 that $A \leq \pi(\mathbf{a}, X)$.

If $\alpha = \beta \cup \gamma$ then since $\tau_w(\beta \cup \gamma) = \tau_w(\beta) \cup \tau_w(\gamma)$ the assumption holds for both β and γ. By the induction hypothesis and closure rule 12, $A \leq \pi(\beta, X) \wedge \pi(\gamma, X)$ so by $\pi 3$, $A \leq \pi(\beta \cup \gamma, X)$.

For the case $\alpha = \beta;\gamma$, note that for all $w = (B, c)$ such that $A \rightarrow^{\beta} B$ all sequences in $\tau_w(\gamma)$ ending in a state $C \leq X$ are green. By closure rule 13 we have $[\beta]\pi(\gamma, X) \in Cl(W)$ and $\pi(\beta, \langle \gamma \rangle X) \in Cl(W)$ and consequently by rules 2 and 3 also $\pi(\gamma, X) \in Cl(W)$ and $\langle \gamma \rangle X \in Cl(W)$. By the induction hypothesis, $B \leq \pi(\gamma, X)$. By Lemma 3, $A \leq [\beta]\pi(\gamma, X)$. Secondly, for all sequences $\sigma_1 : A \rightarrow^{\beta} B \leq \langle \gamma \rangle X$, there exists by Lemma 2 a sequence $\sigma_2 : B \rightarrow^{\gamma} C \leq X$ with $\sigma_1[f] = \sigma_2[i]$. But then $\sigma_1; \sigma_2 \in \tau_w(\beta;\gamma)$ and it follows that σ_1 is green. By the induction hypothesis $A \leq \pi(\beta, \langle \gamma \rangle X)$. It now follows by $\pi 2$ that $A \leq \pi(\beta;\gamma, X)$.

For the case $\alpha = \beta^*$, assume that $A \leq \neg\pi(\beta^*, X)$. By $\pi 6$, $A \leq \langle \beta^* \rangle \neg\pi(\beta, \langle \beta^* \rangle X)$. By closure rule 14, $\langle \beta^* \rangle \neg\pi(\beta, \langle \beta^* \rangle X) \in Cl(W)$. By Lemma 3, there exists a sequence $\sigma_1 : A \rightarrow^{\beta^*} B \leq \neg\pi(\beta, \langle \beta^* \rangle X)$. By closure rule 3, $\neg\pi(\beta, \langle \beta^* \rangle X) \in Cl(W)$, so by the induction hypothesis, there exists a sequence $\sigma_2 : B \rightarrow^{\beta} C \leq \langle \beta^* \rangle X$ with $\sigma_1[f] = \sigma_2[i]$ which is not green. Since $\langle \beta^* \rangle X \in Cl(W)$ by closure rules 2 and 11, we have, from Lemma 3, that there also exists a sequence $\sigma_3 : C \rightarrow^{\beta^*} D \leq X$ with $\sigma_2[f] = \sigma_3[i]$. Thus we have a sequence $\sigma_1; \sigma_2; \sigma_3 : A \rightarrow^{\beta^*} D \leq X$ which is not green, contrary to hypothesis. It follows that $A \leq \pi(\beta^*, X)$.

We are now almost done. After one final lemma, the proof of completeness is trivial.

Lemma 6 For any $X \in Cl(W)$ and state (A, c) of the model, $(A, c) \models X$ iff $A \leq X$.

Proof We use induction on the complexity of X. The basis, $X \in \mathbf{Prop}$, follows directly from the construction of the model. Since for any atom A and any $X \in Cl(W)$, $A \leq X$ if and only if not $A \leq \neg X$, the cases $X = \neg Y$ and $X = Y \vee Z$ are immediate.

For the case $X = \langle \alpha \rangle Y$, note that $w = (A, c) \models \langle \alpha \rangle Y$ iff there exists a sequence $\sigma \in \tau_w(\alpha)$ with $\sigma[f] \models Y$ iff (by induction) there exists an atom $B \leq Y$ with $A \rightarrow^{\alpha} B$ iff (by Lemma 3) $A \leq \langle \alpha \rangle Y$. The arguments for $X = \Diamond(\alpha, Y)$ and $X = \pi(\alpha, Y)$ are similar, using Lemmas 3,4 and 5.

In order to establish completeness it suffices to observe that since $W \in Cl(W)$ is consistent there exists an atom $A \leq W$. By Lemma 6 it follows that for any colour c, $(A, c) \models W$.

6 Conclusion

The logic of this paper has been generalized by Csirmaz [2], who considers models in which transitions which are coloured from arbitrary ordered sets, instead of the set $\{\mathbf{red}, \mathbf{green}\}$. This enables the expression of gradations of deontic modalities, rather than the binary distinction between "permitted" and "not permitted." He also establishes a separation result concerning the interaction of the dynamic modality *after*, which constrains the interconnection structure of models, with the permission modalities, which constrain the way models are coloured.

As mentioned in the introduction, our formulation of the semantics of permission is closely related to the deontic logic of McCarty [8]. McCarty's main semantic construct is the *grand permitted set*, the set of all permitted sequences of states. McCarty uses this to obtain a free-choice permission operator by making an action permitted just when each sequence in its denotation is in the grand permitted set. Further, an action is forbidden when no sequence in its denotation is in the grand permitted set. There are, however, some complications which prevent the straightforward relation of McCarty's logic to DLP by equating the grand permitted set with the set of all green sequences. McCarty's logic is based on an intuitionistic account of states which uses partial states to model actions. He has an interpreted action language in which the primitive action expressions explicitly describe the state changes they produce. This action language is also fully recursive with the deontic language, so that deontic actions such as "making α permitted" may be expressed. Furthermore, observe that in our formulation, any subsequence of a green sequence is green. McCarty, on the other hand, does not assume that any subsequence of a sequence in the grand permitted set is also in the grand permitted set, although he discusses a related condition, closure of the grand permitted set under suffixes. Our completeness result therefore leaves open the problem of completeness of McCarty's logic, although it does cast some light in this direction.

In order to make our logic sufficiently expressive for use in intelligent legal information systems, it will be necessary to add many of the features of McCarty's logic. For any realistic application, it will be necessary to incorporate non-monotonicity and to provide a way to specify deontic actions, i.e. actions which have the effect of changing the set of permitted actions. We have made an initial foray into this area in [9], where we consider how to specify the effects of atomic actions in a logic programming framework. However, much work remains to be done before we have a fully satisfactory treatment of these issues. See [7] for a discussion of the interaction of deontic modalities and non-monotonicity.

Finally, DLP is not yet a complete deontic logic in that it lacks the ability to express *obligation* of actions. It is not difficult to extend the semantics to allow for the inclusion of a notion of obligation. For example, we could introduce an operator $Ob(\alpha)$ with semantics $w \models Ob(\alpha)$ iff all maximal green sequences have a prefix in $\tau_w(\alpha)$. However, it is less clear how to extend the completeness result to incorporate this operator or a generalization, the difficulty being that localization of the model construction to atomic actions cannot straightforwardly be carried out. It appears that something like the expressive power of branching-time temporal logic [3] is necessary to obtain a complete logic containing this modality.

Acknowledgments: This paper represents work done while the author was with the Department of Computer Science, Rutgers University, New Brunswick, NJ. Thanks to Laci Csirmaz and Thorne McCarty for their comments. A version of this paper appeared in the IEEE Symposium on Logic in Computer Science, Philadelphia, June 1990.

References

[1] H.-N. Castaneda. The paradoxes of deontic logic. In R. Hilpinen, editor, *New Studies in Deontic Logic*. Reidel, Dordrecht/Boston, 1981.

[2] Laszlo Csirmaz. Multi-level permission. Technical Report 90-25, DIMACS, Rutgers University, New Brunswick, NJ, April 1990.

[3] E.A. Emerson and A.P. Sistla. Deciding full branching time logic. *Information and Control*, 61(3):175–201, 1984.

[4] M.J. Fisher and R.E. Ladner. Propositional dynamic logic of programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[5] A.J.I. Jones and M. Sergot. Formal specification of security requirements using the theory of normative positions. In *Proc. 2nd European Symposium on Research in Computer Security*, pages 103–121. Springer LNCS No. 648, 1992.

[6] D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.

[7] L. T. McCarty. Defeasible deontic reasoning. *Fundamenta Informaticae*, 21(1):125–148, 1994.

[8] L.T. McCarty. Permissions and obligations. In *Proceedings IJCAI-83*, pages 287–294, August 1983.

[9] R. van der Meyden. A clausal logic for deontic action specification. In V. Saraswat and K. Ueda, editors, *Logic Programming, Proceedings of the 1991 International Symposium*, pages 221–238. MIT Press, 1991.

[10] J.-J. Meyer, H. Weigand, and R. Wieringa. A specification language for static, dynamic and deontic integrity constraints. In J. Demetrovics and B.Thalheim, editors, *Mathematical Fundamentals of Database Systems*. Springer LNCS No. 346, 1989.

[11] J.-J. Ch. Meyer. Free choice permissions and Ross's Paradox: internal vs. external non-determinism. Technical Report IR-130, Dept. Mathematics and Computer Science, Vrije Universiteit, Amsterdam, 1987.

[12] J.-J. Ch. Meyer. A different approach to deontic logic: Deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic*, 29(1), 1988.

[13] N.H. Minsky. The imposition of protocols over open distributed systems. *IEEE Transactions on Software Engineering*, February 1991.

[14] N.H. Minsky. Regularities in software systems. Technical Report LCSR-TR-204, Rutgers University, LCSR, April 1993.

[15] N.H. Minsky and A. Lockman. Ensuring integrity by adding obligations to privileges. In *Proc. 8th International Conference on Software Engineering*, pages 92–102, London, 1985. IEEE.

[16] V. Pratt. Semantical considerations on Floyd-Hoare logic. In *Proceedings 17th IEEE Symposium on Foundations of Computer Science*, pages 109–121, October 1976.

[17] V. Pratt. Process logic. In *Proceedings 6th ACM Symposium on Principles of Programming Languages*, pages 93–100, January 1979.

[18] A. Ross. Imperatives and logic. *Theoria*, 7:53–71, 1941.

[19] D.A. Schlobohm and L.T. McCarty. EPS-II: Estate planning with prototypes. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 1–10. ACM Press, 1989.

[20] K. Segerberg. A completeness theorem in the modal logic of programs. In T. Traczyk, editor, *Universal Algebra and Applications*, volume 9 of *Banach Center Publications*, pages 31–46. PWN Polish Scientific Publishers, Warsaw, 1982.

[21] Y. Shoham. Agent oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.

[22] G. H. von Wright. *An Essay in Deontic Logic and the General Theory of Action.* North Holland, 1968.